arXiv:1411.6463v1 [cs.FL] 24 Nov 2014

Quotient of Acceptance Specifications under Reachability Constraints

Guillaume Verdier and Jean-Baptiste Raclet

IRIT/CNRS, 118 Route de Narbonne, F-31062 Toulouse Cedex 9, France
{verdier,raclet}@irit.fr

Abstract. The quotient operation, which is dual to the composition, is crucial in specification theories as it allows the synthesis of missing specifications and thus enables incremental design. In this paper, we consider a specification theory based on marked acceptance specifications (MAS) which are automata enriched with variability information encoded by acceptance sets and with reachability constraints on states. We define a sound and complete quotient for MAS hence ensuring reachability properties by construction.

1 Introduction

Component-based design aims at building complex reactive systems by assembling components, possibly taken off-the-shelf. This approach can be supported by a specification theory in which requirements correspond to specifications while components are models of the specifications. Such theories come equipped with a set of operations enabling modular system design.

Several recent specification theories are based on modal specifications [16,14,9] including in timed [7,11] or quantitative [1] contexts and with data [2]. In this paper, we introduce *marked acceptance specifications* (MAS): they are based on an extension of modal specifications, called acceptance specifications, which we enrich with marked states to model reachability objectives. This last addition is needed to model session terminations, component checkpoints or rollbacks.

A crucial feature in a specification theory is the operation of quotient. Let S_1 be the specification of a target system and S_2 be the specification of an available black-box component. The specification S_1/S_2 characterizes all the components that, when composed with any model of S_2, conform with S_1. In other words, S_1/S_2 tells what remains to be implemented to realize S_1 while reusing a component doing S_2. By allowing to characterize missing specifications, quotient thus enables incremental design and component reuse.

The quotient of specifications also plays a central role in contract-based design. In essence, a contract describes what a system should guarantee under some assumptions about its context of use. It can be modeled as a pair of specifications (A, G) for, respectively, the assumptions and the guarantees. Satisfiability of a contract then corresponds to the satisfiability of the specification G/A (see [6] for more explanations on contract satisfaction).

Our contribution. Firstly, we define MAS and their semantics. The included marked states allow to specify reachability objectives that must be fulfilled by any model of the MAS. A MAS then characterizes a set of automata called *terminating* as they satisfy the reachability property telling that a marked state can always be reached.

Secondly, we study the compositionality of MAS. We define a compatibility criterion such that two MAS S_1 and S_2 are compatible if and only if the product of any models of S_1 and S_2 is terminating. Further, given two incompatible MAS S_1 and S_2, we propose a construction to refine S_1 into the most general S_1' such that S_1' and S_2 become compatible.

Last, we define the quotient of MAS. This is a two-step construction that makes use of the previous cleaning construction. The operation is shown to be sound and complete.

Related work. Modal specifications [13] enriched with marked states (MMS) have been introduced in [10] for the supervisory control of services. Product of MMS has been investigated in [8]. These papers did not show the need for the more expressive framework of MAS as quotient was not considered. Acceptance specifications have first been proposed in [15] based on [12]. Their non-deterministic version is named Boolean MS in [5]. The LTL model checking of MS has been studied in [4]. However, the reachability considered in this paper can be stated in CTL by `AG(EF(final))` and cannot be captured in LTL.

Quotient of modal and acceptance specifications has been studied in [15,9] and in [3] for the non-deterministic case. It has also been defined for timed [7,11] and quantitative [1] extensions of modal specifications. None of these works consider reachability constraints.

Outline of the paper. We recall some definitions about automata and introduce MAS in Sec. 2. Then, we define the pre-quotient operation in Sec. 3 which only partially solves the problem as it does not ensure the reachability of marked states. In Sec. 4, we give a criterion of compatible reachability telling whether the product of the models of two MAS is always terminating. When this condition of compatible reachability is not met, it is possible to impose some constraints on one of the specifications in order to obtain it, as shown in Sec. 5. Based on this construction, Sec. 6 finally defines the quotient operation on MAS.

2 Modeling with Marked Acceptance Specifications

2.1 Background on automata

A (deterministic) *automaton* over an alphabet Σ is a tuple $M = (R, r^0, \lambda, G)$ where R is a finite set of states, $r^0 \in R$ is the initial state, $\lambda : R \times \Sigma \rightharpoonup R$ is the labeled transition map and $G \subseteq R$ is the set of marked states. The set of *fireable* actions from a state r, denoted $\text{ready}(r)$, is the set of actions a such that $\lambda(r, a)$ is defined.

Given a state r, we define $\text{pre}^*(r)$ and $\text{post}^*(r)$ as the smallest sets such that $r \in \text{pre}^*(r)$, $r \in \text{post}^*(r)$ and for any r', a and r'' such that $\lambda(r', a) = r''$, $r' \in$

$\text{pre}^*(r)$ if $r'' \in \text{pre}^*(r)$ and $r'' \in \text{post}^*(r)$ if $r' \in \text{post}^*(r)$. We also define $\text{pre}^+(r)$ as the union of $\text{pre}^*(r')$ for all r' such that $\exists a : \lambda(r', a) = r$ and $\text{post}^+(r)$ as the union of $\text{post}^*(\lambda(r, a))$ for all $a \in \text{ready}(r)$. Let $\text{Loop}(r) = \text{pre}^+(r) \cap \text{post}^+(r)$.

Two automata M_1 and M_2 are bisimilar iff there exists a simulation relation $\pi : R_1 \times R_2$ such that $(r_1^0, r_2^0) \in \pi$ and for all $(r_1, r_2) \in \pi$, $\text{ready}(r_1) = \text{ready}(r_2) = Z$, $r_1 \in G_1$ iff $r_2 \in G_2$ and for any $a \in Z$, $(\lambda(r_1, a), \lambda(r_2, a)) \in \pi$.

The *product* of two automata M_1 and M_2, denoted $M_1 \times M_2$, is the automaton $(R_1 \times R_2, (r_1^0, r_2^0), \lambda, G_1 \times G_2)$ where $\lambda((r_1, r_2), a)$ is defined as the pair $(\lambda_1(r_1, a), \lambda_2(r_2, a))$ when both $\lambda_1(r_1, a)$ and $\lambda_2(r_2, a)$ are defined.

Given an automaton M and a state r of M, r is a *deadlock* if $r \notin G$ and $\text{ready}(r) = \emptyset$; r belongs to a *livelock* if $\text{Loop}(r) \neq \emptyset$, $G \cap \text{Loop}(r) = \emptyset$ and there is no transition $\lambda(r', a) = r''$ such that $r' \in \text{Loop}(r)$ and $r'' \notin \text{Loop}(r)$. An automaton is *terminating* if it is deadlock-free and livelock-free.

2.2 Marked Acceptance Specification

We now enrich acceptance specifications [15] with marked states to model reachability constraints. The resulting formalism allows to specify a (possibly infinite) set of terminating automata called models.

Definition 1 (MAS). *A* marked acceptance specification *(MAS) over an alphabet Σ is a tuple $S = (Q, q^0, \delta, \text{Acc}, F)$ where Q is a finite set of states, $q^0 \in Q$ is the initial state, $\delta : Q \times \Sigma \rightharpoonup Q$ is the labeled transition map, $\text{Acc} : Q \to 2^{2^\Sigma}$ associates to each state its acceptance set and $F \subseteq Q$ is a set of marked states.*

Basically, an acceptance set is a set of sets of actions a model of the specification is ready to engage in. The *underlying* automaton associated to S is $\text{Un}(S) = (Q, q^0, \delta, F)$. We only consider MAS such that $\text{Un}(S)$ is deterministic.

Definition 2 (Satisfaction). *A terminating automaton M* satisfies *a MAS S, denoted $M \models S$, iff there exists a simulation relation $\pi \subseteq R \times Q$ such that $(r^0, q^0) \in \pi$ and for all $(r, q) \in \pi$:* $\text{ready}(r) \in \text{Acc}(q)$*; if $r \in G$ then $q \in F$; and, for any a and r' such that $\lambda(r, a) = r'$, $(r', \delta(q, a)) \in \pi$. M is called a* model *of S.*

Example 1. A MAS is depicted in Fig. 1(a). Marked states are double-circled while the acceptance sets are indicated near their associated state. The terminating automata M' and M'' in Fig. 1(b) and Fig. 1(c) are models of S_1 because of the respective simulation relation $\pi' = \{(0', 0), (1', 1)\}$ and $\pi'' = \{(0'', 0), (1'', 0), (2'', 1)\}$. Observe that the transitions labeled by b and c are optional in state 0 from the MAS S_1 as these actions are not present in all sets in $\text{Acc}(0)$ and thus may not be present in any model of the specification. Moreover, state 1 in S_1 is marked to encode the constraint that it must be simulated in any model. As a result, although the actions b and c are optional, at least one of the two must be present in any model of S_1. This kind of constraint entails that MAS are more expressive than MS.

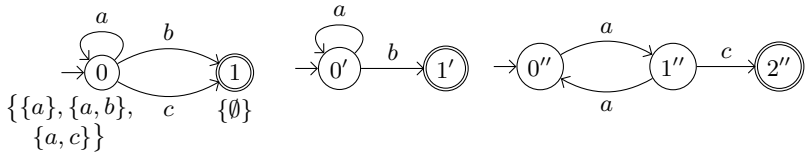

(a) S_1 (b) M' (c) M''

Fig. 1. Example of MAS with two models

The introduced semantic induces some simplifications in the structure of the MAS that we discuss now. This will then lead to the definition of an associated normal form.

- *attractability.* A MAS is said attracted in q when $\text{post}^*(q) \cap F \neq \emptyset$.
- Acc-*consistency.* A state q is Acc-consistent when $\text{Acc}(q) \neq \emptyset$.
- F, Acc-*consistency.* A state q is F, Acc-consistent when $\emptyset \in \text{Acc}(q)$ implies $q \in F$.
- δ, Acc-*consistency.* A state q is δ, Acc-consistent when, for any action $a \in \Sigma$, $\delta(q, a)$ is defined if and only if there exists an $X \in \text{Acc}(q)$ such that $a \in X$.

Remark 1. When $\text{Acc}(q) = \emptyset$, q cannot belong to a simulation relation stating that $M \models S$ as we cannot find an $X \in \text{Acc}(q)$ such that $\text{ready}(r) = X$ for some r. Moreover, when $\emptyset \in \text{Acc}(q)$ and $(r, q) \in \pi$, we can have $\text{ready}(r) = \emptyset$ that is, there is no outgoing transition from r. As M is terminating, this requires that r is marked and thus, q is also marked.

Definition 3 (Normal form). *A MAS is in* normal form *if it is attracted,* Acc-*consistent,* F, Acc-*consistent and* δ, Acc-*consistent in every state* q.

Theorem 1. *Every marked acceptance specification is equivalent to a marked acceptance specification in normal form.*

First, we introduce a specification $S_\perp$ that admits no model. We assume that $S_\perp$ is in normal form. Now proof of the previous theorem is by construction of a MAS in normal form $\rho(S)$ and then proving that S and $\rho(S)$ are equivalent. This construction is detailed in Algo. 1 which returns a MAS corresponding to $\rho(S)$; it defines a pruning operation which removes all the states which are not attracting or not Acc-consistent and updates the acceptance sets and the transition map to enforce F, Acc-consistency and δ, Acc-consistency.

Proof. Following Def. 3, four cases may lead to detect that a MAS S is not in normal form. For each of them, we associate a construction rule to obtain $\rho(S)$ that we analyze now:

- if $\neg$ Acc-consistent(q) then q cannot belong to a simulation relation stating that $M \models S$ as explained in Remark 1. We thus remove the state from S (lines 6 to 9 of Algo. 1);

– if ¬ attracted(q) then no marked state is reachable from q in S. We thus cannot include q in a simulation relation to build a terminating model. As in the previous case, we remove q from S (lines 6 to 9 of Algo. 1);
– if $\neg F$, Acc-consistent(q), we remove $\emptyset$ from Acc(q) (line 12 of Algo. 1). This is a direct consequence of the fact that $\emptyset \in \text{Acc}(q)$ is only relevant when q is marked as advocated in Remark 1;
– last, the combination of items 1 and 3 of Def. 2 indicates that a ready set X is relevant in q if and only if the transition map is defined from q for any $a \in X$. In lines 14 to 17 of Algo. 1, we then update the acceptance sets and the transition map to make them consistent.

As a result, none of these four cases affects the set of models of the received MAS. These construction rules are iteratively applied until a fix-point is reached. The algorithm always finishes as at least the (finite) number of states or the (finite) size of the acceptance sets strictly decreases. □

Algorithm 1 normal_form (S: MAS): MAS

```
1: S' ← S
2: repeat
3:    unchanged ← true
4:    for all q' ∈ Q':
5:       if ¬ attracted(q') ∨ ¬ Acc-consistent(q'):
6:          unchanged ← false
7:          for all q ∈ Q' such that δ'(q, a) = q':
8:             δ'(q, a) = undefined
9:          Q' ← Q' \ {q'}
10:      if ¬F, Acc-consistent(q'):
11:         unchanged ← false
12:         Acc'(q') ← Acc'(q') \ {∅}
13:      if ¬δ, Acc-consistent(q'):
14:         unchanged ← false
15:         for all a ∈ ready'(q') \ ⋃ Acc'(q'):
16:            δ'(q', a) = undefined
17:         Acc'(q') ← {X ∈ Acc'(q') | ∀a ∈ X : δ'(q', a) is defined }
18:      if Q' = ∅:
19:         return S⊥
20: until unchanged
21: return S'
```

As a result of Th. 1, from now on and without loss of generality, we always assume that MAS are in normal form.

At this point, the reader may wonder why attractability and the previous different forms of consistency are not fully part of the definition of MAS. The reason for this is because, in what follows, we propose composition operators on MAS and it is easier to define these constructions without trying to preserve these different requirements. Now if the combination of two MAS (which are now

implicitly supposed to be in normal form) gives rise to a specification violating one of the above requirements then a step of normalization has to be applied on the result in order to have an iterative process.

3 Pre-Quotient Operation of MAS

We first define an operation called *pre-quotient*. Given two MAS S_1 and S_2, it returns a MAS $S_1/\!\!/S_2$ such that the product of any of its models with any model of S_2, *if terminating*, will be a model of S_1. Another operation, defined in Sec. 5, will then be used in Sec. 6 to remove the *"if terminating"* assumption.

Definition 4 (Pre-quotient). *The* pre-quotient *of two MAS S_1 and S_2, denoted $S_1/\!\!/S_2$, is the MAS $(Q_1 \times Q_2, (q_1^0, q_2^0), \delta, \mathrm{Acc}, F)$ with:*

- $\mathrm{Acc}(q_1, q_2) = \{X \mid (\forall X_2 \in \mathrm{Acc}_2(q_2) : X \cap X_2 \in \mathrm{Acc}_1(q_1)) \wedge X \subseteq (\bigcup \mathrm{Acc}_1(q_1)) \cap (\bigcup \mathrm{Acc}_2(q_2))\}$;
- $\forall a \in \Sigma : \delta((q_1, q_2), a)$ *is defined if and only if there exists* $X \in \mathrm{Acc}(q_1, q_2)$ *such that* $a \in X$ *and then* $\delta((q_1, q_2), a) = (\delta_1(q_1, a), \delta_2(q_2, a))$;
- $(q_1, q_2) \in F$ *if and only if* $q_1 \in F_1$ *or* $q_2 \notin F_2$.

Theorem 2 (Correctness). *Given two MAS S_1 and S_2 and an automaton $M \models S_1/\!\!/S_2$, for any $M_2 \models S_2$ such that $M \times M_2$ is terminating, $M \times M_2 \models S_1$.*

Proof. Let $\pi_{/\!/}$ and π_2 be the simulation relations of $M \models S_1/\!\!/S_2$ and $M_2 \models S_2$. Let $\pi \subseteq ((R \times R_2) \times Q_1)$ be the simulation relation such that $((r, r_2), q_1) \in \pi$ if there exists a q_2 such that $(r_2, q_2) \in \pi_2$ and $(r, (q_1, q_2)) \in \pi_{/\!/}$. For any $((r, r_2), q_1) \in \pi$:

- $\mathrm{ready}(r, r_2) \in \mathrm{Acc}_1(q_1)$: by definition of the product of automata, we have: $\mathrm{ready}(r, r_2) = \mathrm{ready}(r) \cap \mathrm{ready}(r_2)$ and by definition of the acceptance set of the pre-quotient, this intersection is in the acceptance set of q_1.
- $(r, r_2) \in G_1 \times G_2$ implies $q_1 \in F_1$: by definition of the pre-quotient, if $r \in G_1$, then $q_1 \in F_1$.
- for any a, if $\lambda((r, r_2), a) = (r', r_2')$, then $\delta_1(q_1, a)$ is defined and moreover, $((r', r_2'), \delta_1(q_1, a)) \in \pi$: $\lambda(r, a) = r'$, so $\delta((q_1, q_2), a)$ is defined and equal to some (q_1', q_2') and, by definition of the pre-quotient, $\delta_1(q_1, a) = q_1'$ and $\delta_2(q_2, a) = q_2'$; $(r, (q_1, q_2)) \in \pi_{/\!/}$ so $(r', (q_1', q_2')) \in \pi_{/\!/}$, $(r_2, q_2) \in \pi_2$, so $(r_2', q_2') \in \pi_2$, hence $((r', r_2'), q_1') \in \pi$. □

The specification returned by the quotient is also expected to be complete, ie., to characterize all the possible automata whose product with a model of S_2 is a model of S_1. However, such a specification may become very large as it will, in particular, have to allow from a state (q_1, q_2) all the transitions which are not fireable from q_2 in S_2. As these transitions will always be removed by the product with models of S_2, they serve no real purpose for the quotient. We propose to return a compact specification for the quotient, without these transitions which we then call *unnecessary*.

An automaton M is said to have no *unnecessary transitions* regarding a MAS S, denoted $M \sim_{\mathcal{U}} S$, if and only if there exists a simulation relation

$\pi \subseteq R \times Q$ such that $(r^0, q^0) \in \pi$ and for all $(r, q) \in \pi$, $\text{ready}(r) \subseteq \bigcup \text{Acc}(q)$ and for every a and r' such that $\lambda(r, a) = r'$, $(r', \delta(q, a)) \in \pi$.

When an automaton M has unnecessary transitions regarding a MAS S, it is possible to remove these transitions. Let $\rho_u(M, S)$ be the automaton $M' = (R \times Q, (r^0, q^0), \lambda', G \times Q)$ with:

$$\lambda'((r, q), a) = \begin{cases} (\lambda(r, a), \delta(q, a)) \text{ if } a \in \bigcup \text{Acc}(q) \\ \text{undefined otherwise} \end{cases}$$

Theorem 3. *Given an automaton M and a MAS S, we have: $\rho_u(M, S) \sim_{\mathcal{U}} S$. Moreover, for all $M_S \models S$, the automata $M \times M_S$ and $\rho_u(M, S) \times M_S$ are bisimilar.*

Proof. $\rho_u(M, S) \sim_{\mathcal{U}} S$: let π be the simulation relation such that for any state (r, q) of $\rho_u(M, S)$, $((r, q), q) \in \pi$; by definition of ρ_u, $\text{ready}(r, q) \subseteq \bigcup \text{Acc}(q)$.

$M \times M_S$ and $\rho_u(M, S) \times M_S$ are bisimilar: let π be the simulation relation such that for any state (r, r_S) of $M \times M_S$ and $((r, q), r_S)$ of $\rho_u(M, S) \times M_S$, $((r, r_S), ((r, q), r_S)) \in \pi$. Then, $\text{ready}(((r, q), r_S)) = (\text{ready}(r) \cap \bigcup \text{Acc}(q)) \cap \text{ready}(r_S)$. As r_S implements q, $\text{ready}(r_S) \subseteq \bigcup \text{Acc}(q)$, so $\text{ready}(((r, q), r_S)) = \text{ready}(r) \cap \text{ready}(r_S) = \text{ready}((r, r_S))$. □

We can then prove that our pre-quotient is complete for automata without unnecessary transitions. Given an arbitrary automaton, it suffices to remove these transitions with ρ_u before checking if it is a model of the quotient.

Theorem 4. *Given two MAS S_1 and S_2 and an automaton M such that $M \sim_{\mathcal{U}} S_2$ and for all $M_2 \models S_2$ we have $M \times M_2 \models S_1$, then $M \models S_1 /\!\!/ S_2$.*

Proof. Let π be a simulation relation such that $(r^0, (q_1^0, q_2^0)) \in \pi$ and for any $(r, (q_1, q_2)) \in \pi$, a and r' such that $\lambda(r, a) = r'$, $(r', \delta((q_1, q_2), a)) \in \pi$. This definition of π is only correct if for any $(r, (q_1, q_2)) \in \pi$ and a such that $\lambda(r, a)$ is defined, $\delta((q_1, q_2), a) = (\delta_1(q_1, a), \delta_2(q_2, a))$ is also defined. As $M \sim_{\mathcal{U}} S_2$, $a \in \bigcup \text{Acc}_2(q_2)$, so there exists an $X \in \text{Acc}_2(q_2)$ such that $a \in X$ and then $\delta_2(q_2, a)$ is defined (as S_2 is well-formed). As $\delta_2(q_2, a)$ is defined, there exists an automaton $M_2 \models S_2$ with a state r_2 implementing q_2 such that (r, r_2) is reachable in $M \times M_2$ and $\lambda_2(r_2, a)$ is defined. Then, $\lambda((r, r_2), a)$ is defined and, as $M \times M_2 \models S_1$, it implies that $\delta((q_1, q_2), a)$ is defined.

There are then three points to prove for any $(r, (q_1, q_2)) \in \pi$:

- $\text{ready}(r) \in \text{Acc}((q_1, q_2))$: by definition of the pre-quotient, $\text{ready}(r)$ must verify two properties:
  - $\forall X_2 \in \text{Acc}_2(q_2) : \text{ready}(r) \cap X_2 \in \text{Acc}_1(q_1)$:
    Let X_2 be an element of $\text{Acc}_2(q_2)$. There exists an automaton M_2 with a state r_2 such that (r, r_2) is reachable in $M \times M_2$ and $\text{ready}(r_2) = X_2$. Then, as $M \times M_2 \models S_1$ by a simulation relation $\pi_\times$ and $((r, r_2), q_1) \in \pi_\times$, $\text{ready}(r) \cap \text{ready}(r_2) = \text{ready}(r) \cap X_2 \in \text{Acc}_1(q_1)$.

- $\text{ready}(r) \subseteq \bigcup \text{Acc}_1(q_1) \cap \bigcup \text{Acc}_2(q_2)$:
  By definition of $\sim_{\mathcal{U}}$, $\text{ready}(r) \subseteq \bigcup \text{Acc}_2(q_2)$.
  Assume that $\text{ready}(r) \not\subseteq \bigcup \text{Acc}_1(q_1)$: there is an $a \in \text{ready}(r)$ such that $a \notin \bigcup \text{Acc}_1(q_1)$. As M has no unnecessary transition regarding S_2, there is a model M_2 of S_2 with a state r_2 such that (r, r_2) is reachable in $M \times M_2$ and $a \in \text{ready}(r_2)$. Then, the transition $((r, r_2), a)$ is defined in $M \times M_2$. As $M \times M_2 \models S_1$, the transition (q_1, a) has to be defined, which is in contradiction with the hypothesis that $a \notin \bigcup \text{Acc}_1(q_1)$. Thus, $\text{ready}(r) \subseteq \bigcup \text{Acc}_1(q_1)$.

– $r \in G$ implies $(q_1, q_2) \in F_{/\!/}$, that is $q_1 \in F_1$ or $q_2 \notin F_2$:
  This property is only false if $r \in G$, $q_1 \notin F_1$ and $q_2 \in F_2$. In this case, there exists an automaton $M_2 \models S_2$ with a state r_2 such that (r, r_2) is reachable in $M \times M_2$ and $r_2 \in G_2$. Then, $M \times M_2 \models S_1$ by a simulation relation $\pi_\times$, $((r, r_2), q_1) \in \pi_\times$ and (r, r_2) is marked. By definition of satisfaction, it implies that $q_1 \in F_1$, which is impossible as we already know that $q_1 \notin F_1$. So $r \in G$ implies $(q_1, q_2) \in F_{/\!/}$.
– for any a and r' such that $\lambda(r, a) = r'$, $(r', \delta((q_1, q_2), a)) \in \pi$ is trivial by definition of π. □

Corollary 1 (Completeness). *Given two MAS S_1 and S_2 and an automaton M such that for all $M_2 \models S_2$ we have $M \times M_2 \models S_1$, then $\rho_u(M, S_2) \models S_1 /\!/ S_2$.*

Proof. By Th. 3, we know that $\rho_u(M, S_2) \sim_{\mathcal{U}} S_2$ and for any $M_2 \models S_2$, $\rho_u(M, S_2) \times M_2$ is bisimilar to $M \times M_2$, which implies that $\rho_u(M, S_2) \times M_2 \models S_1$. Then, Th. 4 implies that $\rho_u(M, S_2) \models S_1 /\!/ S_2$. □

Observe now that the pre-quotient $S_1 /\!/ S_2$ may admit some models whose product with some models of S_2 may not be terminating. Consider indeed the specifications S_1 and S_2 of Fig. 1(a) and 2(a) and their pre-quotient in Fig. 2(b). The product of the models M_1^1 of $S_1 /\!/ S_2$ (Fig. 2(c)) and M' of S_2 (Fig. 1(b)) has a livelock and thus is not terminating. One may think that there is an error in the pre-quotient computation and that it should not allow to realize only $\{a, c\}$, without b in $\text{Acc}(0, 0')$. Indeed, it would forbid the model M_1^1, but it would also disallow some valid models such as M_1^2 (Fig. 2(d)), which realizes $\{a, c\}$ in a state and $\{a, b\}$ in another, thus synchronizing on b with any model of S_2 and allowing the joint reachability of the marked states.

In the next section, we define a criterion allowing to test whether the product of any models of two MAS is terminating or not. On this basis, we will then refine the pre-quotient in Sec. 6 in order to guarantee the reachability property.

4 Compatible Reachability of MAS

By definition, the product of some models of two MAS may not terminate due to two different causes, namely deadlock and livelock. We consider separately the two issues to derive a compatible reachability criterion on MAS.

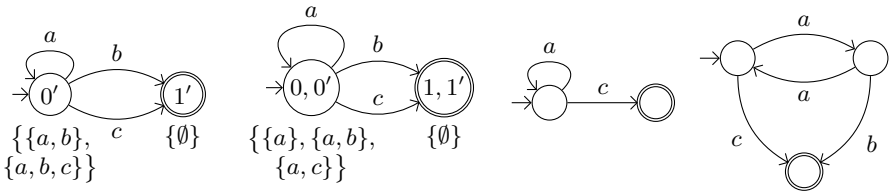

(a) S_2 (b) $S_1 /\!\!/ S_2$ (c) $M_1^1 \models S_1 /\!\!/ S_2$ (d) $M_1^2 \models S_1 /\!\!/ S_2$

Fig. 2. Example of pre-quotient

4.1 Deadlock-free specifications

In this section, we propose a test to check if two MAS S_1 and S_2 have some models M_1 and M_2 such that $M_1 \times M_2$ has a deadlock. To do so, we characterize deadlock-free pairs of states, from which no deadlock may arise in the product of any two models of S_1 and S_2.

Given two acceptance sets A_1 and A_2, let $\text{Compat}(A_1, A_2)$ be true iff for all $X_1 \in A_1$ and $X_2 \in A_2$, $X_1 \cap X_2 \neq \emptyset$. Now a pair of states (q_1, q_2) is said to be deadlock-free, denoted $\text{DeadFree}(q_1, q_2)$, if $\text{Acc}_1(q_1) = \text{Acc}_2(q_2) = \{\emptyset\}$ or $\text{Compat}(\text{Acc}_1(q_1), \text{Acc}_2(q_2))$.

Definition 5 (Deadlock-free MAS). *Two MAS S_1 and S_2 are* deadlock-free *when all the reachable pairs of states in $\text{Un}(S_1) \times \text{Un}(S_2)$ are deadlock-free.*

Theorem 5. *Two MAS S_1 and S_2 are deadlock-free if and only if for any $M_1 \models S_1$ and $M_2 \models S_2$, $M_1 \times M_2$ is deadlock-free.*

Proof. ($\Rightarrow$) Suppose that (r_1, r_2) is a deadlock in $M_1 \times M_2$. Then (r_1, r_2) is not marked and $\text{ready}((r_1, r_2)) = \emptyset$. Now $\text{ready}((r_1, r_2)) = \text{ready}(r_1) \cap \text{ready}(r_2)$ and moreover, $(r_1, q_1) \in \pi_1$ and $(r_2, q_2) \in \pi_2$ implies $\text{ready}(r_1) \in \text{Acc}_1(q_1)$ and $\text{ready}(r_2) \in \text{Acc}_2(q_2)$. As a result, for $X_1 = \text{ready}(r_1) \in \text{Acc}_1(q_1)$, $X_2 = \text{ready}(r_2) \in \text{Acc}_2(q_2)$, we have: $X_1 \cap X_2 = \emptyset$ and thus $\neg\,\text{Compat}(\text{Acc}_1(q_1), \text{Acc}_2(q_2))$. Moreover, (r_1, r_2) is not marked so (q_1, q_2) is not marked and $\emptyset \notin \text{Acc}_1(q_1)$ and $\emptyset \notin \text{Acc}_2(q_2)$. In consequence, we have $\neg\,\text{DeadFree}(q_1, q_2)$ and S_1 and S_2 are not deadlock-free.

($\Leftarrow$) Suppose that S_1 and S_2 are not deadlock-free: there exists q_1 and q_2 such that $\neg\,\text{DeadFree}(q_1, q_2)$. Then there exists $X_1 \in \text{Acc}_1(q_1)$ and $X_2 \in \text{Acc}_2(q_2)$ which verify $X_1 \cap X_2 = \emptyset$. For any $M_1 \models S_1$ and $M_2 \models S_2$ with $(r_1, q_1) \in \pi_1$ and $(r_2, q_2) \in \pi_2$ such that $\text{ready}(r_1) = X_1$ and $\text{ready}(r_2) = X_2$, we have $\text{ready}((r_1, r_2)) = X_1 \cap X_2 = \emptyset$ in $M_1 \times M_2$. Moreover, $\text{Acc}_1(q_1) \neq \{\emptyset\}$ (or $\text{Acc}_2(q_2) \neq \{\emptyset\}$), so there exists a model of S_1 (resp. S_2) such that a state r implementing q_1 (resp. q_2) is not marked and has at least one transition leading to another marked state, so (r_1, r_2) is not marked. As a result, (r_1, r_2) is a deadlock and $M_1 \times M_2$ is not deadlock-free. □

4.2 Livelock-free specifications

In this section, we explain how we can check if two MAS S_1 and S_2 have some models M_1 and M_2 such that $M_1 \times M_2$ has a livelock. We identify the cycles shared between S_1 and S_2 along with the transitions leaving them. We check if at least one of these transitions is preserved in the product of any two models of S_1 and S_2. Before studying these common cycles, a first step consists in unfolding S_1 and S_2 so as possible synchronizations become unambiguous.

Unfolding. Given two specifications S_1 and S_2, we define the *partners* of a state q_1 as $Q_2(q_1) = \{q_2 \mid (q_1, q_2)$ is reachable in $\mathrm{Un}(S_1) \times \mathrm{Un}(S_2)\}$; the set $Q_1(q_2)$ is defined symmetrically. As a shorthand, if we know that a state q_1 has exactly one partner, we will also use $Q_2(q_1)$ to denote this partner.

If some states of S_2 have several partners, it is possible to transform S_2 so that each of its states has at most one partner, while preserving the set of models of the specification. The unfolding of S_2 in relation to S_1 is the specification $((Q_1 \cup \{q^?\}) \times Q_2, (q_1^0, q_2^0), \delta_u, \mathrm{Acc}_u, (Q_1 \cup \{q^?\}) \times F_2)$ where:

- $q^?$ is a fresh state ($q_1^?$ denotes a state in $Q_1 \cup \{q^?\}$);
- $\delta_u((q_1^?, q_2), a)$ is defined if and only if $\delta_2(q_2, a)$ is defined and then:

$$\delta_u((q_1, q_2), a) = \begin{cases} (\delta_1(q_1, a), \delta_2(q_2, a)) & \text{if } \delta_1(q_1, a) \text{ is defined} \\ (q^?, \delta_2(q_2, a)) & \text{otherwise} \end{cases}$$
$$\delta_u((q^?, q_2), a) = (q^?, \delta_2(q_2, a))$$

- $\mathrm{Acc}_u((q_1^?, q_2)) = \mathrm{Acc}_2(q_2)$.

Two MAS S_1 and S_2 have *single partners* if and only if for all $q_1 \in Q_1$, we have $|Q_2(q_1)| \leq 1$ and for all $q_2 \in Q_2$, we also have $|Q_1(q_2)| \leq 1$.

Given two MAS S_1 and S_2, there exists some MAS S_1' and S_2', called *unfoldings* of S_1 and S_2, with single partners and which have the same models as S_1 and S_2. These two MAS can be computed by unfolding S_1 in relation to S_2 and then S_2 in relation to the unfolding of S_1.

Lemma 1. *Given two MAS S_1 and S_2 and S_u the unfolding of S_2 in relation to S_1, $S_u \equiv S_2$.*

Proof. ($\Rightarrow$) Let M be a model of S_u. Let π_u be the simulation relation between the states of M and the states of S_u and let π_2 be the simulation relation such that $(r, q_2) \in \pi_2$ if and only if there exists a $q_1^?$ such that $(r, (q_1^?, q_2)) \in \pi_u$. $(r^0, q_2^0) \in \pi_2$ and for any $(r, q_2) \in \pi_2$:

- $\mathrm{ready}(r) \in \mathrm{Acc}_2(q_2)$ as $\mathrm{ready}(r) \in \mathrm{Acc}_u((q_1^?, q_2)) = \mathrm{Acc}_2(q_2)$;
- if $r \in G$, $q_2 \in F_2$ as $(q_1^?, q_2) \in (Q_1 \cup \{q^?\}) \times F_2$;
- for any $a \in \mathrm{ready}(r)$, $(\lambda(r, a), \delta_2(q_2, a)) \in \pi_2$ as $(\lambda(r, a), \delta_u((q_1^?, q_2), a)) \in \pi_u$.

Thus M is a model of S_2.

($\Leftarrow$) Let M be a model of S_2. Let π_2 be the simulation relation between the states of M and the states of S_2 and let π_u be the simulation relation such that $(r, (q_1^?, q_2)) \in \pi_u$ if and only if $(r, q_2) \in \pi_2$ and $(q_1^?, q_2)$ is reachable in S_u. $(r^0, (q_1^0, q_2^0)) \in \pi_u$ and for any $(r, (q_1^?, q_2)) \in \pi_2$:

– $\text{ready}(r) \in \text{Acc}_u((q_1^?, q_2))$ as $\text{ready}(r) \in \text{Acc}_2(q_2) = \text{Acc}_u((q_1^?, q_2))$;
– if $r \in G$, $(q_1^?, q_2) \in (Q_1 \cup \{q^?\}) \times F_2$ as $q_2 \in F_2$;
– for any $a \in \text{ready}(r)$, $(\lambda(r, a), \delta_u((q_1^?, q_2), a)) \in \pi_u$ as $(\lambda(r, a), \delta_2(q_2, a)) \in \pi_2$.
Thus M is a model of S_u. □

Lemma 2. *Given two MAS S_1 and S_2 and S_u the unfolding of S_2 in relation to S_1, for any $(q_1, (q_1^?, q_2))$ reachable in $\text{Un}(S_1) \times \text{Un}(S_2)$, $q_1 = q_1^?$.*

Proof. If a state is reachable in $\text{Un}(S_1) \times \text{Un}(S_2)$, there is a path from the initial state to it. By induction on this path:

– if it is empty, we are in the initial state $(q_1^0, (q_1^0, q_2^0))$;
– otherwise, we are in a state $(q_1, (q_1, q_2))$ and there is a transition by an action a to another state $(\delta_1(q_1, a), \delta_u((q_1, q_2), a))$. As $\delta_1(q_1, a)$ is defined, $\delta_u((q_1, q_2), a) = (\delta_1(q_1, a), \delta_2(q_2, a))$, so the destination state is the pair $(\delta_1(q_1, a), (\delta_1(q_1, a), \delta_2(q_2, a)))$. □

Lemma 3. *Given two MAS S_1 and S_2 and S_u the unfolding of S_2 in relation to S_1, for any state q_u of S_u, $|Q_1(q_u)| \leq 1$.*

Proof. Let suppose that $|Q_1(q_u)| > 1$. Then, there exists at least two different states q_1 and q_1' such that (q_1, q_u) and (q_1', q_u) are reachable in $\text{Un}(S_1) \times \text{Un}(S_2)$. By the definition of the unfolding operation, there exists some $q_1^? \in Q_1 \cup \{q_?\}$ and $q_2 \in Q_2$ such that $q_u = (q_1^?, q_2)$. By Lem. 2, $q_1 = q_1^?$ and $q_1' = q_1^?$, so $q_1 = q_1'$. But we know by hypothesis that they are different, so $|Q_1(q_u)| \leq 1$. □

Two MAS S_1 and S_2 have *single partners* if and only if for all $q_1 \in Q_1$, we have $|Q_2(q_1)| \leq 1$ and for all $q_2 \in Q_2$, we also have $|Q_1(q_2)| \leq 1$.

Theorem 6. *Given two MAS S_1 and S_2, there exists some MAS S_1' and S_2', called* unfoldings *of S_1 and S_2, with single partners and which are equivalent to S_1 and S_2.*

Proof. Let S_1' be the unfolding of S_1 in relation to S_2 and S_2' the unfolding of S_2 in relation to S_1'.

By Lem. 1, we know that S_1' has the same models as S_1 and S_2' as S_2.

By Lem. 3, we know that for any q_1' in S_1', $|Q_2(q_1')| \leq 1$ and that for any q_2' in S_2', $|Q_1'(q_2')| \leq 1$. Remains to prove that $|Q_2'(q_1')| \leq 1$.

Let q_1' be a state of S_1'. If $|Q_2(q_1')| = 0$, then $|Q_2'(q_1')| = 0$ as S_2 and S_2' have the same models. Otherwise, there exists a q_2 such that $Q_2(q_1') = \{q_2\}$. There exists then n states (with $n > 0$) q_{2_i}' of the form $(q_{1_i}^?, q_2)$. But each q_{2_i}' is in relation with at most one state (q_{1_i}') of S_1', as $|Q_1'(q_{2_i}')| \leq 1$, and all these q_{1_i}' are different (as the q_{2_i}' are different). So there is at most one q_{2_i}' in relation with q_1' and thus $|Q_2'(q_1')| \leq 1$. □

Cycles. In order to detect livelocks, we need to study the cycles that may be present in the models of a specification. Intuitively, a cycle is characterized by its states and the transitions between them.

Given a MAS S, the partial map $C : Q \rightharpoonup 2^{\Sigma}$ represents a *cycle* in S if and only if for any $q \in \mathrm{dom}(C)$, (a) $C(q) \neq \emptyset$, (b) $\exists X \in \mathrm{Acc}(q)$ such that $C(q) \subseteq X$, (c) $\mathrm{dom}(C) \subseteq \mathrm{post}^*(q)$ and (d) $\forall a \in C(q) : \delta(q, a) \in \mathrm{dom}(C)$.

A model M of a MAS S *implements* a cycle C if and only if there exists a set $\mathcal{R}$ of states of M such that each $q \in \mathrm{dom}(C)$ is implemented by at least one state of $\mathcal{R}$ and for each $r \in \mathcal{R}$ and for each q such that $(r, q) \in \pi$, (a) $q \in \mathrm{dom}(C)$, (b) $C(q) \subseteq \mathrm{ready}(r)$, (c) $\forall a \in C(q) : \lambda(r, a) \in \mathcal{R}$ and (d) $\forall a \in \mathrm{ready}(r) \backslash C(q) : \lambda(r, a) \notin \mathcal{R}$. A cycle is said to be *implementable* if there exists a model M of S implementing the cycle.

We define in Algo. 2 an operation, $\mathrm{Loop}_{\models}$-rec, which computes the cycles of a MAS passing by a given state.

Theorem 7. *Given a MAS S, a model M of S implements a cycle $\mathcal{C}$ if and only if $\mathcal{C} \in \mathrm{Loop}_{\models}(S)$.*

Proof. ($\Rightarrow$) Let $\mathcal{C}$ be a cycle in S and M a model of S implementing $\mathcal{C}$, with $\mathcal{R}$ the set of states of M implementing the states of $\mathcal{C}$. Let r be an element of $\mathcal{R}$ and q a state it implements. By definition, $\mathrm{Loop}_{\models}(S)$ contains the result of $\mathrm{Loop}_{\models}(S, q)$, which calls $\mathrm{Loop}_{\models}\text{-rec}(S, q, \emptyset)$. For an iteration of the loop at line 5, the variable C will take the value of $\mathcal{C}(q)$ and it will be inserted in the generated cycle. The algorithm will then be called recursively on the successors of q in the cycle, until q is reached again, thus obtaining $\mathcal{C}$.

($\Leftarrow$) Let $\mathcal{C}$ be a cycle returned by $\mathrm{Loop}_{\models}(S)$. It is possible to build an automaton M implementing the states and transitions of $\mathcal{C}$. The problem is to make sure that this automaton is terminating, ie. that it is possible to reach a marked state from any implementation of a state of $\mathrm{dom}(\mathcal{C})$. By definition of the cycle implementation relation, we know that $\mathrm{dom}(\mathcal{C}) \cap F \neq \emptyset$ or $\exists q \in \mathrm{dom}(\mathcal{C}) : \exists X \in \mathrm{Acc}(q) : \mathcal{C}(q) \subset X$. In the first case, there is a marked state in the loop, thus M is terminating. In the second case, we know that there is a state q, implemented in M by a state r, from which there is a transition by an action a which leaves the cycle, that is, $\lambda(r, a)$ is not in the set of states implementing $\mathcal{C}$. There is thus no constraint on the transitions from $\lambda(r, a)$ and it will be possible to reach a marked state from it (provided that S is well-formed). So M is terminating and in consequence is a model of S. □

However, some of these cycles may not be implementable. For instance, the cycle $\mathcal{C} = \{0 \mapsto \{a\}\}$ is not implementable in the MAS depicted in Figure 3, as all the models have to eventually realize the transition by b to reach the marked state and then are not allowed to simultaneously realize the transition by a.

In order to be implementable, a cycle has to contain a marked state or it must be possible to realize a transition that is not part of the cycle in addition to the transitions of the cycle. Thus, the set of implementable cycles of a MAS S, denoted $\mathrm{Loop}_{\models}(S)$, is $\bigcup_{q \in Q} \{\mathcal{C} \in \mathrm{Loop}_{\models}\text{-rec}(S, q, \emptyset) \mid \mathrm{dom}(\mathcal{C}) \cap F \neq \emptyset \vee \exists q_{\mathcal{C}} \in \mathrm{dom}(\mathcal{C}) : \exists X \in \mathrm{Acc}(q_{\mathcal{C}}) : \mathcal{C}(q_{\mathcal{C}}) \subset X\}$.

Livelock-freeness. We can now analyze the cycles of two MAS with single partners in order to detect if there may be a livelock in the product of some of

Fig. 3. A MAS over {a,b} with no implementable cycle

Algorithm 2 $\text{Loop}_{\models}$-rec (S: MAS, q: State, cycle: Cycle): Set Cycle

```
1: if q ∈ dom(cycle) then return {cycle}
2: res ← ∅
3: for all A ∈ Acc(q):
4:    cycle_acc ← {a | a ∈ A ∧ q ∈ post*(δ(q, a))}
5:    for all C ∈ 2^cycle_acc \ {∅}:
6:       current ← {cycle}
7:       for all a ∈ C:
8:          current ← ⋃_{cycle∈current} Loop_⊨-rec(S, δ(q, a), cycle ∪ {q ↦ C})
9:       res ← res ∪ current
10: return res
```

their models. To do so, we distinguish two kinds of transitions: those, denoted $\mathcal{A}$, which are always realized when the cycle is implemented and those, denoted $\mathcal{O}$, which may (or may not) be realized when the cycle is implemented. These sets are represented by partial functions from a state to a set of sets of actions and given, for a particular cycle $\mathcal{C}$, by the following formulae:

$$\mathcal{A} = \{q \mapsto \text{leaving}(q, A) \mid A \notin \text{Acc}(q), (q, A) \in \mathcal{C}\}$$
$$\mathcal{O} = \{q \mapsto \text{leaving}(q, A) \mid A \in \text{Acc}(q), (q, A) \in \mathcal{C} \wedge \text{leaving}(q, A) \neq \emptyset\}$$
$$\text{where leaving}(q, A) = \{X \setminus A \mid X \in \text{Acc}(q) \wedge A \subset X\}$$

Definition 6. *Given two MAS S_1 and S_2 with single partners and a cycle $\mathcal{C}_1$ in S_1 such that all its states have a partner, $\mathcal{C}_1$ is livelock-free in relation to S_2, denoted* $\text{LiveFree}(\mathcal{C}_1, S_2)$*, if and only if, when the cycle $\mathcal{C}_2 = \{Q_2(q) \mapsto \mathcal{C}_1(q) \mid q \in \text{dom}(\mathcal{C}_1)\}$ is in $\text{Loop}_{\models}(S_2)$:*

1. *$\mathcal{A}_{\mathcal{C}_1} \neq \emptyset$, $\mathcal{A}_{\mathcal{C}_2} \neq \emptyset$ and there exists $q'_1 \in \text{dom}(\mathcal{A}_{\mathcal{C}_1})$ such that $Q_2(q'_1) \in \text{dom}(\mathcal{A}_{\mathcal{C}_2})$ and $\text{Compat}(\mathcal{A}_{\mathcal{C}_1}(q'_1), \mathcal{A}_{\mathcal{C}_2}(Q_2(q'_1)))$, or*
2. *$\mathcal{A}_{\mathcal{C}_1} \neq \emptyset$, $\mathcal{A}_{\mathcal{C}_2} = \emptyset$, $\text{dom}(\mathcal{C}_2) \cap F_2 = \emptyset$ and $\forall q'_2 \in \text{dom}(\mathcal{O}_{\mathcal{C}_2}) : Q_1(q'_2) \in \text{dom}(\mathcal{A}_{\mathcal{C}_1})$ and $\text{Compat}(\mathcal{A}_{\mathcal{C}_1}(Q_1(q'_2)), \mathcal{O}_{\mathcal{C}_2}(q'_2))$, or*
3. *$\mathcal{A}_{\mathcal{C}_1} = \emptyset$, $\mathcal{A}_{\mathcal{C}_2} \neq \emptyset$, $\text{dom}(\mathcal{C}_1) \cap F_1 = \emptyset$ and $\forall q'_1 \in \text{dom}(\mathcal{O}_{\mathcal{C}_1}) : Q_2(q'_1) \in \text{dom}(\mathcal{A}_{\mathcal{C}_2})$ and $\text{Compat}(\mathcal{O}_{\mathcal{C}_1}(q'_1), \mathcal{A}_{\mathcal{C}_2}(Q_2(q'_1)))$.*

Definition 7 (Livelock-free specifications). *Two MAS S_1 and S_2 with single partners are* livelock-free *if all the implementable cycles of S_1 are livelock-free in relation to S_2.*

This definition only tests the implementable cycles of S_1. It is not necessary to do the symmetrical test (checking that the implementable cycles of S_2 verify

LiveFree) because we only compare the cycle of S_1 with the same cycle in S_2 and the tests of Def. 6 are symmetric.

The previous definition offers a necessary and sufficient condition to identify MAS which can have two respective models whose product has a livelock:

Theorem 8. *Two MAS S_1 and S_2 with single partners are livelock-free if and only if for any $M_1 \models S_1$ and $M_2 \models S_2$, $M_1 \times M_2$ is livelock-free.*

Proof. ($\Rightarrow$) Assume that there exists $M_1 \models S_1$, $M_2 \models S_2$ such that $M_1 \times M_2$ has a livelock, that is, there exists (r_1, r_2) such that $\text{Loop}((r_1, r_2)) \neq \emptyset$, $\text{Loop}((r_1, r_2)) \cap G = \emptyset$ and there is no transition (r', a, r'') such that $r' \in \text{Loop}((r_1, r_2))$ and $r'' \notin \text{Loop}((r_1, r_2))$.

- If there exists a cycle $\mathcal{C}_1 \in \text{Loop}_{\models}(S_1)$ which is implemented in M_1 by the states of $\text{Loop}(r_1)$ and $\mathcal{C}_2 = \{Q_2(q) \mapsto \mathcal{C}_1(q) \mid q \in \text{dom}(\mathcal{C}_1)\}$ is implemented in M_2 by the states of $\text{Loop}(r_2)$:
  - if there is no transition leaving $\text{Loop}(r_1)$, then $\mathcal{A}_{\mathcal{C}_1} = \emptyset$ and $\text{dom}(\mathcal{C}_1) \cap F_1 \neq \emptyset$, so the three tests of Def. 6 fail and S_1 and S_2 are not livelock-free; symmetrically S_1 and S_2 are not livelock-free if there is no transition leaving $\text{Loop}(r_2)$;
  - if there are transitions leaving $\text{Loop}(r_1)$ and $\text{Loop}(r_2)$, they are not compatible, ie. they have different actions or different source states. If in both models, some of these transitions are in $\mathcal{A}$ (they have to be present whenever the cycle is implemented), the test 1 of Def. 6 will detect that they are not compatible. If there are some transitions in $\mathcal{A}_{\mathcal{C}_1}$ but none in $\mathcal{A}_{\mathcal{C}_2}$, test 2 will detect that M_2 may implement a transition that will not be covered by the transitions in $\mathcal{A}_{\mathcal{C}_1}$. Test 3 handles the symmetrical case. Finally, if there are transitions neither in $\mathcal{A}_{\mathcal{C}_1}$ nor $\mathcal{A}_{\mathcal{C}_2}$, it is always possible to generate a livelock and all three tests fail.
- Otherwise, multiple cycles are implemented simultaneously in the model by unfolding them or two slightly different cycles are implemented in M_1 and M_2, and then there will also be a livelock in the models which implement only one of the cycles, which brings us back to the first case.

($\Leftarrow$) Assume that S_1 and S_2 are not livelock-free. Then, there exists a cycle $\mathcal{C}_1$ such that $\neg\,\text{LiveFree}(\mathcal{C}_1, S_2)$. Then, the three conditions of Def. 7 are all false.

- If $\mathcal{A}_{\mathcal{C}_1} \neq \emptyset$ and $\mathcal{A}_{\mathcal{C}_2} \neq \emptyset$, then for any state $q'_1 \in \text{dom}(\mathcal{A}_{\mathcal{C}_1})$ in S_1, we have $\neg\,\text{Compat}(\mathcal{A}_{\mathcal{C}_1}(q'_1), \mathcal{A}_{\mathcal{C}_2}(Q_2(q'_1)))$. So there exists a model M_1 of S_1 implementing $\mathcal{C}_1$ and a model M_2 of S_2 implementing $\mathcal{C}_2$ such that there is no transition leaving the cycle in their product, hence there is a livelock in $M_1 \times M_2$.
- If $\mathcal{A}_{\mathcal{C}_1} \neq \emptyset$ and $\mathcal{A}_{\mathcal{C}_2} = \emptyset$, there exists a state $q'_2 \in \text{dom}(\mathcal{O}_{\mathcal{C}_2})$ such that we have $\neg\,\text{Compat}(\mathcal{A}_{\mathcal{C}_1}(Q_1(q'_2)), \mathcal{O}_{\mathcal{C}_2}(q'_2))$. So for any model M_1 of S_1 implementing $\mathcal{C}_1$, its product with a model M_2 of S_2 implementing $\mathcal{C}_2$ for which the only transition leaving the cycle is from an implementation of q'_2 will have a livelock.
- If $\mathcal{A}_{\mathcal{C}_1} = \emptyset$ and $\mathcal{A}_{\mathcal{C}_2} \neq \emptyset$, we are in the case symmetric to the previous one.

– If $\mathcal{A}_{\mathcal{C}_1} = \emptyset$ and $\mathcal{A}_{\mathcal{C}_2} = \emptyset$, either one of the specifications has no transitions leaving the cycle ($\mathcal{O}_{\mathcal{C}_i} = \emptyset$ too), so there are some models such that their product has a livelock, or both $\mathcal{O}_{\mathcal{C}_1}$ and $\mathcal{O}_{\mathcal{C}_2}$ are not empty, and then there exists an $M_1 \models S_1$ implementing $\mathcal{C}_1$ such that the only transition(s) leaving the cycle is (are) from a state r_1 and an $M_2 \models S_2$ implementing $\mathcal{C}_2$ such that the only transition(s) leaving the cycle is (are) from a state r_2 which is never paired with r_1 in $M_1 \times M_2$, hence there is a livelock in $M_1 \times M_2$. □

Specifications with compatible reachability. By combining the tests for deadlock-free and livelock-free specifications, we can define a criterion checking if two MAS S_1 and S_2 have some models M_1 and M_2 such that $M_1 \times M_2$ is not terminating.

Definition 8 (Compatible reachability). *Two MAS S_1 and S_2 have a* compatible reachability*, denoted $S_1 \sim_{\mathcal{T}} S_2$, if and only if they are deadlock-free and their unfoldings are livelock-free. They have an incompatible reachability otherwise.*

Theorem 9. *Given two MAS S_1 and S_2, $S_1 \sim_{\mathcal{T}} S_2$ if and only if for any $M_1 \models S_1$ and $M_2 \models S_2$, $M_1 \times M_2$ is terminating.*

Proof. By definition, $S_1 \sim_{\mathcal{T}} S_2$ if and only if S_1 and S_2 are deadlock-free and livelock-free. By Th. 5, 6 and 8, this is equivalent to: for any $M_1 \models S_1$ and $M_2 \models S_2$, $M_1 \times M_2$ is deadlock-free and livelock-free, that is, $M_1 \times M_2$ is terminating. □

This theorem allows *independent implementability* of MAS: given two MAS with compatible reachability, each specification may be implemented independently from the other while keeping the guarantee that the composition of the resulting implementations will be terminating and thus satisfy by construction a reachability property.

5 Correction of MAS with Incompatible Reachability

We now define an operation that, given two MAS S_1 and S_2 with incompatible reachability, returns a MAS refining S_1 with a compatible reachability with S_2.

5.1 Deadlock correction

First, given two non-deadlock-free MAS S_1 and S_2, we propose to refine S_1 such that the obtained MAS S_1' is deadlock-free with S_2. For this, we iterate through all the non-deadlock-free pairs of states (q_1, q_2) and remove the elements of the acceptance set of q_1 which may cause a deadlock, as described in Algo. 3. Note that it may return an empty specification, because of ρ, which then means that for any model M_1 of S_1, there exists a model M_2 of S_2 such that $M_1 \times M_2$ has a deadlock.

Theorem 10 (Deadlock correction). *Given two MAS S_1 and S_2, $M_1 \models S_1$ is such that for any $M_2 \models S_2$, $M_1 \times M_2$ is deadlock-free if and only if $M_1 \models$ dead_correction(S_1, S_2).*

Proof. ($\Rightarrow$) Assume that for any $M_1 \models S_1$ and $M_2 \models S_2$, $M_1 \times M_2$ is deadlock-free. By Th. 5, S_1 and S_2 are deadlock-free, which implies that there is no pair of states (q_1, q_2) such that $\neg\,\mathrm{DeadFree}(q_1, q_2)$. Thus, the set dead_pairs in Algo. 3 is empty and $\mathrm{dead_correction}(S_1, S_2) = S_1$, so $M_1 \models \mathrm{dead_correction}(S_1, S_2)$.

($\Leftarrow$) Assume that there exists an $M_2 \models S_2$ such that $M_1 \times M_2$ has a deadlock pair of states (r_1, r_2). By Th. 5, this implies that S_1 and S_2 are not deadlock-free and thus that there exists a pair of states (q_1, q_2) (implemented by (r_1, r_2)) reachable in $\mathrm{Un}(S_1) \times \mathrm{Un}(S_2)$ such that $\neg\,\mathrm{DeadFree}(q_1, q_2)$. Then, in $\mathrm{dead_correction}(S_1, S_2)$, either the acceptance set of q_1 has been reduced so that $\mathrm{Compat}(\mathrm{Acc}'_1(q_1), \mathrm{Acc}_2(q_2))$ is true and $\mathrm{DeadFree}(q_1, q_2)$ or q_1 is not reachable anymore and then (q_1, q_2) is not reachable in $\mathrm{Un}(S_1) \times \mathrm{Un}(\mathrm{dead_correction}(S_1, S_2))$. Consequently, either $\mathrm{ready}(r_1) \notin \mathrm{Acc}'_1(q_1)$ or $(r_1, q_1) \notin \pi$, and thus M_1 is not a model of $\mathrm{dead_correction}(S_1, S_2)$. □

Algorithm 3 dead_correction (S_1: MAS, S_2: MAS): MAS

1: $S'_1 \leftarrow S_1$
2: **for all** (q_1, q_2) such that $\neg\,\mathrm{DeadFree}(q_1, q_2)$:
3: **if** $\mathrm{Acc}_2(q_2) = \{\emptyset\}$:
4: **if** $\emptyset \in \mathrm{Acc}'_1(q_1)$ **then** $\mathrm{Acc}'_1(q_1) \leftarrow \{\emptyset\}$ **else** $\mathrm{Acc}'_1(q_1) \leftarrow \emptyset$
5: **else**:
6: $\mathrm{Acc}'_1(q_1) \leftarrow \{X_1 \mid X_1 \in \mathrm{Acc}'_1(q_1) \wedge \forall X_2 \in \mathrm{Acc}_2(q_2) : X_1 \cap X_2 \neq \emptyset\}$
7: **return** $\rho(S'_1)$

5.2 Livelock correction

Secondly, given S_1 and S_2 two deadlock-free MAS, we propose to refine S_1 such that the obtained specification S'_1 is livelock-free with S_2.

There are two ways to prevent livelocks from occuring in the product of the models of two MAS: removing some transitions so that states from which it is not possible to guarantee termination will not be reached and forcing some transitions to be eventually realized in order to guarantee that it will be possible to leave cycles without marked states. For this last method, we introduce marked acceptance specifications with priorities that are MAS with some priority transitions which have to be eventually realized.

Definition 9 (MAS with priorities). *A marked acceptance specification with priorities (MASp) is a tuple $(Q, q^0, \delta, \mathrm{Acc}, P, F)$ where $(Q, q^0, \delta, \mathrm{Acc}, F)$ is a MAS and $P : 2^{2^{Q \times \Sigma}}$ is a set of* priorities.

Definition 10 (Satisfaction). *An automaton M implements a MASp S if M implements the underlying MAS and for all $\mathcal{P} \in P$, either $\forall (q,a) \in \mathcal{P} : \forall r : (r,q) \notin \pi$ or $\exists (q,a) \in \mathcal{P} : \exists r : (r,q) \in \pi \wedge a \in \text{ready}(r)$.*

Intuitively, P represents a conjunction of disjunctions of transitions: at least one transition from each element of P must be implemented by the models of the specification.

Let S_1 and S_2 be two MAS and q_1 a state of S_1 such that q_1 belongs to a livelock. Then, there exists a cycle $\mathcal{C}_1$ in S_1 and its partner $\mathcal{C}_2$ in S_2 such that the conditions given in Def. 6 are false. Given these cycles, Algo. 4 ensures that the possible livelock will not happen, either by adding some priorities or by removing some transitions.

Algorithm 4 live_corr_cycle (S_1: MASp, $\mathcal{C}_1$: Cycle, S_2: MAS, $\mathcal{C}_2$: Cycle): MASp

if $\mathcal{A}_{\mathcal{C}_2} \neq \emptyset$:
  $Q_A \leftarrow \{q_1 \mid Q_2(q_1) \in \text{dom}(\mathcal{A}_{\mathcal{C}_2}) \wedge \forall A \in \mathcal{A}_{\mathcal{C}_2}(Q_2(q_1)) : A \cap \text{ready}(q_1) \neq \emptyset\}$
  if $Q_A \neq \emptyset$:
    $P \leftarrow \{\bigcup_{1 \leq i \leq |Q_A|} \{(q_i, a) \mid a \in X_i\} \mid X_i \in \{A \cap \text{ready}(q_i) \mid A \in \mathcal{A}_{\mathcal{C}_2}(Q_2(q_i))\}\}$
    return $(Q_1, q_1^0, \delta_1, \text{Acc}_1, P_1 \cup P, F_1)$
else if $\text{dom}(\mathcal{C}_2) \cap F_2 = \emptyset$:
  $\text{Acc}' \leftarrow \text{Acc}_1$
  for all $q_1 \in \{Q_1(q_2) \mid q_2 \in \text{dom}(\mathcal{O}_{\mathcal{C}_2})\}$:
    $\text{Acc}'(q_1) \leftarrow \{X \mid X \in \text{Acc}_1(q_1) \wedge \forall O \in \mathcal{O}_{\mathcal{C}_2}(Q_2(q_1)) : X \cap O \neq \emptyset\}$
  return $\rho((Q_1, q_1^0, \delta_1, \text{Acc}', P_1, F_1))$
$\text{Acc}' \leftarrow \text{Acc}_1$
for all $q_1 \in Q_1$:
  $\text{Acc}'(q_1) \leftarrow \{X \mid X \in \text{Acc}_1(q_1) \wedge \forall a \in X : \delta(q_1, a) \notin \text{dom}(\mathcal{C}_1)\}$
return $\rho((Q_1, q_1^0, \delta_1, \text{Acc}', P_1, F_1))$

We then iterate over the possible cycles, fixing those which may cause a livelock , as described in Algo. 5.

Algorithm 5 live_correction (S_1: MAS, S_2: MAS): MASp

$S_1' \leftarrow (Q_1, q_1^0, \delta_1, \text{Acc}_1, \emptyset, F_1)$
for all $\mathcal{C}_1 \in \text{Loop}_{\models}(S_1)$ s.t. $\forall q_1 \in \text{dom}(\mathcal{C}_1) : |Q_2(q_1)| = 1$ and $\neg\,\text{LiveFree}(\mathcal{C}_1, S_2)$:
  $S_1' \leftarrow$ live_corr_cycle$(S_1', \mathcal{C}_1, S_2, \{Q_2(q) \mapsto \mathcal{C}_1(q) \mid q \in \text{dom}(\mathcal{C}_1)\})$
return S_1'

Theorem 11 (Livelock correction). *Given two MAS S_1 and S_2, $M_1 \models S_1$ is such that for any $M_2 \models S_2$, $M_1 \times M_2$ is livelock-free if and only if $M_1 \models$ live_correction(S_1, S_2).*

Proof. ($\Rightarrow$) Assume that for any $M_1 \models S_1$ and $M_2 \models S_2$, $M_1 \times M_2$ is livelock-free. By Th. 8, S_1 and S_2 are livelock-free which means, by Def. 7, that for any implementable cycle $\mathcal{C}_1$ in S_1 such that its states have a partner in S_2, we have LiveFree($\mathcal{C}_1, S_2$). In this case, the test at line 2 of Algo. 5 is always false and so live_correction(S_1, S_2) returns S_1, of which M_1 is a model by hypothesis.

($\Leftarrow$) Assume that there exists an $M_2 \models S_2$ such that $M_1 \times M_2$ has a livelock.

- If there exists a cycle $\mathcal{C}_1 \in \text{Loop}_{\models}(S_1)$ which is implemented in M_1 by the states of the loop in which there is a livelock when combined with M_2. Thus, live_correction_cycle will be called with $\mathcal{C}_1$. There are three cases:
  - If $\mathcal{A}_{\mathcal{C}_2}$ is not empty, some transitions are present in all the models of S_2 implementing $\mathcal{C}_2$, so the models of S_1 should realize (at least) one of these transitions once. If it is possible, some priorities are added, see lines 3 to 5 of Algo. 4. This addition will only remove the models of S_1 that never realize any transition in $\mathcal{A}_{\mathcal{C}_2}$ and thus that will have a livelock with some models of M_2 (which only realize the transition of $\mathcal{A}_{\mathcal{C}_2}$.
  - If $\mathcal{A}_{\mathcal{C}_2}$ is empty but there is no marked state in $\mathcal{C}_2$, all the models of S_2 implementing $\mathcal{C}_2$ will eventually realize a transition of $\mathcal{O}_{\mathcal{C}_2}$ in order to reach a marked state (as there is none in the cycle). The only way to avoid a livelock with any model of S_2 is to realize all the transitions that these models may use to reach a marked state, which is done in lines 6 to 10.
  - Otherwise, there will always be a possible livelock with some models of S_2, so the only possibility is to disallow all the models which implement this cycle, which is done in lines 10 to 14.

  So M_1 is not a model of the MASp returned by live_correction_cycle for $\mathcal{C}_1$ and thus it is not a model of live_correction(S_1, S_2).
- Otherwise, multiple cycles are implemented simultaneously and there will also be livelocks in the models which implement only one of the cycles. As argued in the previous item, applying live_correction_cycle for these cycles will generate a specification forbidding the corresponding models, and then M_2 will not be a model of the resulting specification as it only combines the behavior of these models. □

As a result, by applying successively dead_correction and live_correction, we can define the following operation $\rho_{\mathcal{T}}$:

$$\rho_{\mathcal{T}}(S_1, S_2) = \text{live_correction}(\text{dead_correction}(S_1, S_2), S_2)$$

Given two MAS S_1 and S_2, it refines the set of models of S_1 as precisely as possible so that their product with any model of S_2 is terminating.

Theorem 12 (Incompatible reachability correction). *Given two MAS S_1 and S_2, $M \models \rho_{\mathcal{T}}(S_1, S_2)$ if and only if $M \models S_1$ and for any $M_2 \models S_2$, $M \times M_2$ is terminating.*

Proof. For any $M \models \rho_{\mathcal{T}}(S_1, S2)$ and $M_2 \models S_2$, $M \times M_2$ is terminating if and only if $M \times M_2$ is deadlock-free and livelock-free. By theorems 5 and 8, this is

true if and only if $\rho_{\mathcal{T}}(S_1, S_2)$ and S_2 are deadlock-free and livelock-free, which is true by definition of $\rho_{\mathcal{T}}$ and Theorems 10 and 11. □

6 Quotient Operation of MAS

We can now combine the pre-quotient and cleaning operations to define the quotient of two MAS.

Definition 11. *Given two MAS S_1 and S_2, their* quotient S_1/S_2 *is given by* $\rho_{\mathcal{T}}(S_1 /\!\!/ S_2, S_2)$.

Theorem 13 (Soundness). *Given two MAS S_1 and S_2 and an automaton $M \models S_1/S_2$, for any $M_2 \models S_2$, $M \times M_2 \models S_1$.*

Proof. By Th. 12, we know that for any $M_2 \models S_2$, $M \times M_2$ is terminating. Thus, Th. 2 implies that $M \times M_2 \models S_1$. □

Theorem 14 (Completeness). *Given two MAS S_1 and S_2 and an automaton M such that $\forall M_2 \models S_2 : M \times M_2 \models S_1$, then $\rho_u(M, S_2) \models S_1/S_2$.*

Proof. We know by Corollary 1 that $\rho_u(M, S_2) \models S_1 /\!\!/ S_2$. We then deduce by Th. 12 that $\rho_u(M, S_2) \models S_1/S_2$. □

These theorems indicate that each specification S_2 and S_1/S_2 may be implemented independently from the other and that the composition of the resulting implementations will eventually be terminating and will also satisfy S_1.

7 Conclusion

In this paper, we have introduced marked acceptance specifications. We have developed several compositionality results ensuring a reachability property by construction and, in particular, a sound and complete quotient. Note that this framework can almost immediately be enriched with a refinement relation, parallel product and conjunction by exploiting the constructions available in [15] and [8], hence providing a complete specification theory as advocated in [16].

Considering an acceptance setting instead of a modal one offers a gain in terms of expressivity as MAS provide more flexibility than the marked extension of modal specifications [8]. This benefit becomes essential for the quotient as may/must modalities are not rich enough to allow for a complete operation. Consider indeed the two MMS S_1 and S_2 in Figure 4 in which optional transitions are represented with dashed arrows while required transitions are plain arrows. A correct and complete modal quotient in this example would tell in the initial state of S_1/S_2 that at least one action between a and b is required. This cannot be encoded by modalities but it can be correctly stated by the acceptance set $\{\{a\}, \{b\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}\}$.

Observe also that quotient of two MAS is heterogeneous in the sense that its result may be a MASp. By definition, MASp explicitly require to *eventually*

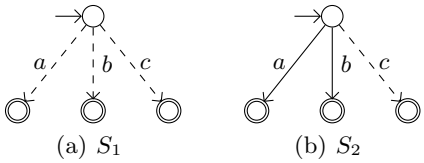

(a) S_1 (b) S_2

Fig. 4. Two MMS showing that a modal quotient cannot exist

realize some transitions fixed in the priority set P. By bounding the delay before the implementation of the transitions, a MASp could become a standard MAS and the quotient would then become homogeneous. Algorithms for bounding MASp are left for future investigations.

References

1. Bauer, S.S., Fahrenberg, U., Juhl, L., Larsen, K.G., Legay, A., Thrane, C.R.: Weighted modal transition systems. Formal Methods in System Design 42(2), 193–220 (2013)
2. Bauer, S.S., Larsen, K.G., Legay, A., Nyman, U., Wasowski, A.: A modal specification theory for components with data. Sci. Comput. Program. 83, 106–128 (2014)
3. Benes, B., Delahaye, B., Fahrenberg, U., Kretínský, J., Legay, A.: Hennessy-milner logic with greatest fixed points as a complete behavioural specification theory. In: CONCUR. LNCS, vol. 8052, pp. 76–90. Springer (2013)
4. Benes, N., Cerná, I., Kretínský, J.: Modal transition systems: Composition and LTL model checking. In: ATVA. LNCS, vol. 6996, pp. 228–242. Springer (2011)
5. Benes, N., Kretínský, J., Larsen, K.G., Møller, M.H., Srba, J.: Parametric modal transition systems. In: ATVA. LNCS, vol. 6996, pp. 275–289. Springer (2011)
6. Benveniste, A., Raclet, J.B., Caillaud, B., Nickovic, D., Passerone, R., Sangiovanni-Vincentelli, A., A., H.T., Larsen, K.G.: Contracts for the design of embedded systems - part II: theory (2012)
7. Bertrand, N., Legay, A., Pinchinat, S., Raclet, J.B.: Modal event-clock specifications for timed component-based design. Sci. Comput. Program. 77(12), 1212–1234 (2012)
8. Caillaud, B., Raclet, J.B.: Ensuring reachability by design. In: ICTAC. LNCS, vol. 7521, pp. 213–227. Springer (2012)
9. Chen, T., Chilton, C., Jonsson, B., Kwiatkowska, M.Z.: A compositional specification theory for component behaviours. In: ESOP. LNCS, vol. 7211, pp. 148–168. Springer (2012)
10. Darondeau, P., Dubreil, J., Marchand, H.: Supervisory control for modal specifications of services. In: WODES. pp. 428–435 (2010)
11. David, A., Larsen, K.G., Legay, A., Nyman, U., Wasowski, A.: Timed I/O automata : A complete specification theory for real-time systems. In: HSCC. pp. 91–100. ACM (2010)
12. Hennessy, M.: Acceptance trees. J. ACM 32(4), 896–928 (1985)
13. Larsen, K.G., Thomsen, B.: A modal process logic. In: LICS. pp. 203–210. IEEE (1988)

14. Lüttgen, G., Vogler, W.: Modal interface automata. Logical Methods in Computer Science 9(3) (2013)
15. Raclet, J.B.: Residual for component specifications. In: FACS. Electr. Notes Theor. Comput. Sci., vol. 215, pp. 93–110 (2008)
16. Raclet, J.B., Badouel, E., Benveniste, A., Caillaud, B., Legay, A., Passerone, R.: A modal interface theory for component-based design. Fun. Informaticae 108(1-2), 119–149 (2011)